This work has been submitted to the IEEE for possible publication. Copyright may be transferred without notice, after which this version may no longer be accessible

arXiv:1612.06702v1 [cs.RO] 20 Dec 2016

Efficient Optical flow and Stereo Vision for Velocity Estimation and Obstacle Avoidance on an Autonomous Pocket Drone

Kimberly McGuire[1], Guido de Croon[1], Christophe De Wagter[1], Karl Tuyls[2] and Hilbert Kappen[3]

***Abstract*—Micro Aerial Vehicles (MAV) are very suitable for flying in indoor environments, but autonomous navigation is challenging due to their strict hardware limitations. This paper presents a highly efficient computer vision algorithm called Edge-FS for the determination of velocity and depth. It runs at 20 Hz on a 4 g stereo camera with an embedded STM32F4 microprocessor (168 MHz, 192 kB) and uses feature histograms to calculate optical flow and stereo disparity. The stereo-based distance estimates are used to scale the optical flow in order to retrieve the drone's velocity. The velocity and depth measurements are used for fully autonomous flight of a 40 g pocket drone only relying on on-board sensors. The method allows the MAV to control its velocity and avoid obstacles.**

I. Introduction

Deployment of Micro Aerial Vehicles (MAV) is valuable for indoor tasks such as search-and-rescue operations, green house observations and more. Tiny MAVs, also called pocket drones(<50 grams, as in Fig. 1), are ideal for maneuvering through very narrow spaces, as often occurs in indoor environments. In order for them to autonomously navigate through a GPS-deprived area, there are several on-board sensors to consider (laser rangers, motion sensors, infrared, sonar). The pocket drone's sensor of choice is a light weight camera. It is the most energy efficient and versatile option, as multiple variables can be observed from a single image stream: distance, motion, object recognition and more.

A camera enables the MAV to extract essential information needed for autonomous navigation, for instance stereo vision for obstacle avoidance [1]. Since there are strict limitations on energy expenditure, sensing, and processing capabilities on a pocket drone, even relatively efficient stereo vision methods [2][3] are computationally too heavy to run on-board a microprocessor. Therefore, an even more efficient stereo vision algorithm was developed, which is able to run at 10 Hz on a 20 g flapping wing MAV, the DelFly Explorer [4]. It is still the lightest fully autonomous MAV to this date, which can fly through a room and avoid obstacles with purely onboard sensing and processing [5].

Since tailed flapping wing MAVs such as the DelFly Explorer are passively stable, there is no need to compute their velocity to compensate for drift. However, for inherently unstable platforms like a quadcopter, velocity estimation is necessary for stabilization. Optical flow is the way in which objects move in an image and is the most important visual cue for velocity estimation. It can be calculated in a dense manner (Horn-Schunck[6], Färneback[7]) or a sparse manner, e.g., by tracking features such as Shi-Tomasi[8] or FAST[9] over time with a Lucas-Kanade tracker. These types of techniques have proven themselves on numerous occasions [10], nonetheless do require a platform with a decent amount of computing power. On a pocket drone such standard optical flow methods either cannot be run in real-time or will occupy a large percentage of its processing power. Especially when autonomous flight is the final goal, this is considered a waste of resources, as optical flow determination will only constitute a part of what the MAV has to do.

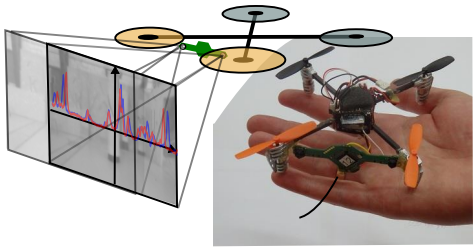

Fig. 1. Pocket drone with a light with a forward looking 4 g stereo camera. A very efficient vision algorithm runs embedded on the STM32F4 processor (168 MHz, 192 kB), to determine velocity and depth necessary for the pocket drone's visual navigation.

To this end, we have drawn inspiration from the study in [11], which proposed using feature-histograms to track motion in the image. Specifically, in [12], we presented an efficient algorithm to estimate optical flow. This method, named EdgeFlow, improved upon the work in [11] by introducing a variable time horizon for determining sub-pixel flow. EdgeFlow ran embedded at 20 Hz on a light weight stereo camera positioned underneath a pocket drone. The stereo camera was facing down and was estimating optical flow and a global height estimate. With these, the MAV determined its own velocity and used this in a guided control, where it autonomously matched externally-given velocity references. However, a 4 g stereo camera for a 40 g pocket drone is significant, so it is a waste to have this "heavy" sensor looking downward and and not using it to avoid obstacles in the flight direction.

This paper presents a major extension of EdgeFlow, which enables the stereo camera to face forward on a MAV, so it can be used for navigation purposes. As the pocket drone will now be facing hallways, rooms, doors etc., the assumption of looking straight at a flat plane will not hold anymore. The same matching paradigm used to determine EdgeFlow, will now be used to not only calculate optical flow but also stereo depth over the entire image. *EdgeStereo*, as called for convenience,

[1] Faculty of Aerospace Engineering, Delft University of Technology, NL k.n.mcguire@tudelft.nl [2] Department of Computer Science, University of Liverpool, UK [3] Faculty of Science, Radboud University of Nijmegen, NL

uses the so-determined distances to properly scale the locally observed optical flow in order to retrieve a velocity estimate. This combination of EdgeFlow and EdgeStereo will be called *Edge-FS*.

Our main contribution is that the presented method provides both velocity and distance estimates, while still being computationally efficient enough to run close to the frame rate on a very limited embedded processor. As such, the method enables unstable MAVs such as tiny quadcopters to perform fully autonomous flights in unknown environments. The EdgeFlow and EdgeStereo methods will be explained in more detail in section II. Off-line results for velocity estimates with a set of images is shown in section II-C. From here, the algorithm is embedded on the lightweight stereo camera and placed on 40 gr pocket drone for velocity estimation (section III-B). Finally, the velocity estimate is used together with EdgeStereo-based obstacle detection to perform fully autonomous navigation in an environment with obstacles (section III-C). This is followed by some concluding remarks.

A. Related Work

In related research, several works have achieved optical flow based control of a MAV ([13][14][15]). As mentioned in the introduction, the standard optical flow methods are computationally too heavy to run on a quadcopter of less than 50 grams. For instance, Dunkley et al. have flown with a 25 gram quadcopter before, while computing optical flow for visual odometry[16]. However, this was done on an external computer. As miniaturization of hardware also poses a limitation on communication bandwidth, this can result in a significant delay in the controls. To obtain full autonomy, it would be wise to uncouple a MAV of any external dependencies.

To design extremly lightweight MAVs for autonomous flight, some researchers looked into EMD sensors [17] and other 1D signal sensors[18]. Briod et al.[19] proposed the design of a 45 gram quadcopter for optical flow based control with 1D flow sensors. They followed up with this research on a heavier 278 gram platform containing 8 of these sensors [20], pointing in all directions. With this they could hover the quadcopter in different cluttered environments. The results are impressive, nevertheless results were achieved by using multiple single purpose sensors. As they can only sense motion, it does not leave much room to detect other variables necessary for navigation.

More similar to our research, Moore et al. implemented an efficient optic flow algorithm on a small lightweight (2gr) omnidirectional camera system on a 30 gram helicopter [21]. With a ring of 8 low-resolution image chips (64 x 64 pixels), the MAV could compute optical flow. It did this by computing the edges, compressing the images and calculate the displacement by block matching which resulted in translational optical flow. The vision calculations where done on-board the helicopter with 10 Hz, yet the flight controls where computed off-board. Although the potential of a full on-board implementation is there, the redundancy lies in the ratio of cameras to sensed variables. One camera has the potential of detecting flow in 3 directions; they used 8 to only detect 2 (forward and sideways velocity).

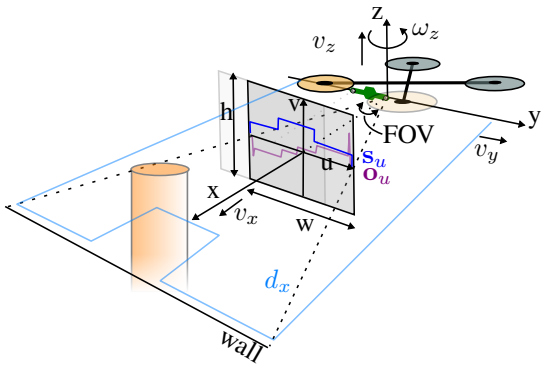

Fig. 2. The MAV's body fixed coordinates with respects to the camera axis, shown for the left camera (XYZ). The image coordinates in width and height are represented as u and v respectively.

Oleynikova et al. developed a reactive controller for a slightly larger quadcopter (30 cm in diameter) for obstacle avoidance, based on depth information [22]. From the obtained stereo disparity map, they accumulated the values along the columns to get a summed disparity factor. Assuming that the obstacles are vertical and long, these can be detected quickly. The stereo map was calculated over the entire image first before accumulation to a vector. This significantly impacts the amount of computation making it less suitable for implementation on a smaller MAV.

II. Velocity and Depth from Edges

To achieve autonomous navigation with a camera on an unstable pocket drone, we need to obtain two variables: velocity and depth. In the introduction we mentioned that many of the main stream computer vision will be computationally too heavy to run on the pocket drone. In [12], we presented EdgeFlow, which can detect optical flow within the image in a semi-dense but computationally efficient manner, embedded on a 4 g stereo board. During the experiments, the stereo camera looked down to the ground, estimating the pocket drone's forward and sideways velocity. This section will explain the modifications that are necessary to make the stereo camera point forward and still be able to measure those variables. EdgeFlow will be concisely recapped. Subsequently, we will present its extension with EdgeStereo to Edge-FS, which will be used for obstacle detection in the experiment part of this paper.

A. From Camera to State

When looking orthogonally at a planar ground surface while moving, the optical flow field, is rather simple and allows for easy determination of the forward and sideways velocities with the help of a single height measurement. To not bump into anything in the process, one needs to face forwards, but then the optical flow field will become more complex. This is due to the likely non-planar (3D structure) of the environment. Moreover, the forward velocity now can only be observed by

means of the divergence of optical flow, which is more difficult to determine, especially close to the focus of expansion. Below, we delve into how we determine the velocities with the help of forward images. In principle, the unscaled velocities and the rotation rates can be determined from the image alone, according to the paper of [23]. The paper of Longuet et al. implies that measured flow ($\mathbf{o}_u$) is the summation of a translational flow ($\mathbf{o}_u^T$) and rotational component ($\mathbf{o}_u^R$). Before estimating the horizontal planar velocity, we first have to determine $\mathbf{o}_u^R$.

Although [23] assumes rotations in all directions, we can make simpler assumptions for the pocket drone. To fly to whatever the camera points at is essential for obstacle avoidance. Here, correctional pitch and roll motion for drift compensation will be relatively small, but yaw rotations will be more common. Assuming that the latter only has significant effect on the optical flow, $\mathbf{o}_u^R$ can be approximated (assuming small angles) using the gyroscopes on the on-board IMU of the pocket drone:

$$\mathbf{o}_u^R \approx \omega_Z \cdot \frac{w}{FOV} \tag{1}$$

where w is the width of the image, FOV is the field of view and ω_Z is the yaw rotation measured from the gyroscopes.

Now that $\mathbf{o}_u^R$ is known, we can isolate $\mathbf{o}_u^T$ to determine the pocket drone's forward (v_x) and sideways velocity. With the coordinate system we use in this paper (Fig. 2), Longuet's equation of $\mathbf{o}_u^T$ is expressed as:

$$\mathbf{o}_u^T = (-v_y + \mathbf{x}v_x)/\mathbf{p}_x \tag{2}$$

$$\mathbf{d}_x\mathbf{o}_u^T = (-v_y + \mathbf{x}v_x) \tag{3}$$

Depth, $\mathbf{d}_x$, scales the optical flow resulting in motion parallax, as close objects appear to move faster than objects far away. In [12], a global height estimate was used to scale the optical flow back to velocities, which is sufficient if the camera is looking at a flat floor or perpendicular to a straight wall. This assumption will not hold when the MAV is flying towards a wall at an angle, or through a hallway with doors. This non-constant depth needs to be accounted for when scaling the optical flow, therefore the stereo depth is needed over the entire size of the image for a better velocity estimate. Local right-left image disparity from a stereo camera can be transformed to actual depth in meters by using the camera parameters, with the following approximation:

$$\mathbf{d}_x \approx \frac{w \cdot r}{(FOV \cdot \mathbf{s}_u)} \tag{4}$$

where r is the distance between the two cameras, and the stereo disparity in pixels is $\mathbf{s}_u$.

With depth $\mathbf{d}_x$ and translational optical flow $\mathbf{o}_u^T$, it is now possible to calculate the MAV's sideways and forward velocity by fitting a linear model to (3). Next we will explain how to obtain both optical flow and stereo depth from a stream of low resolution stereo images.

B. Procedure for Edge-FS

The matching principle of EdgeFlow is shown in Fig. 3. The images A and B are first filtered with (a) a Sobel filter to get the horizontal gradients. These gradients are summed along the rows (compressed) to (b) an edge histogram. From image A and B, the edge histograms are compared with (c) Sum of Absolute Differences (SAD) block matching. This locates the similar patches of the edge histograms within a certain distance from each other, to obtain the pixel displacement between the two.

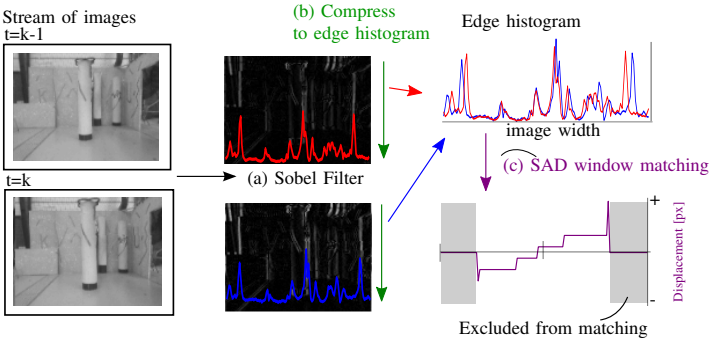

Fig. 3. The matching algorithm for both EdgeFlow as EdgeStereo. The images' gradients (a) calculated a Sobel filter, (b) summed up to an edge histogram. These are (c) matched with other edge histogram. The gray areas are excluded sections (equal to the range plus half SAD block size).

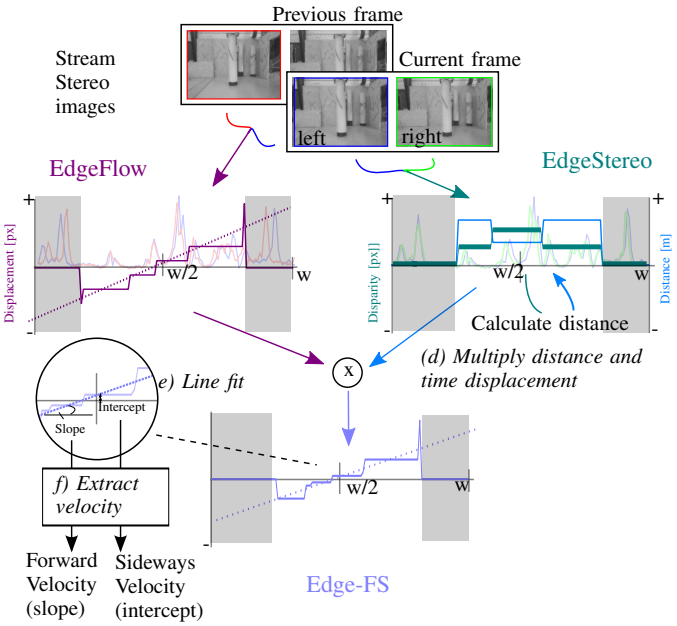

Fig. 4. The temporal pixel disparities per column of EdgeFlow is (d) scaled by EdgeStereo. Following (3), e) a line fit is done on these values, from which the forward and sideways velocities can be extracted from the slope and intercept, respectively.

If image A and B from Fig. 3 are two temporal sequential images (t with $t-1$) in time, this will result in the pixel flow, thus EdgeFlow as presented in [12]. Based on the previous flow value, EdgeFlow adaptively chooses how far in time ($t-n$) it will compare the current edge-histogram to. On top of that, the flow shift predicted by a yaw rotation ($\mathbf{o}_u^R$, as calculated in (1) will shift the start of the block matching scheme. This and the adaptive time horizon will be present for the experiments in this paper. Note that in [12], also the direction along the image height was used to estimate the

forward velocity (as the camera was looking down). In this paper, it will not be used as that value (v_x) will now be subtracted from the divergence of Edge-FS.

Previously in [12], the entire edge histogram of the left and right image was matched to obtain a global depth estimate. To get a better velocity estimate with a forward camera, we need to use pixel disparity per column. To calculate both column-wise optical flow and stereo vision and keep the algorithm computationally efficient, the exact same matching principle of EdgeFlow (Fig. 3) is used, resulting in EdgeStereo. Disparity to depth in meters is calculated with the known camera parameters and (4) from the last section. (d) EdgeStereo scales EdgeFlow to compensate for the motion parallax (see Fig. 4), which results in the left side of (3). These values will then (e) be fitted to a linear model, which gives us the slope and intercept of the line. With the camera parameters, (f) the forward and sideways velocities are estimated.

C. Off-line Vision Experiments

Before implementing the algorithm on the actual stereo board, EdgeFlow was run on a set of stereo-images in MATLAB (version R2015b on a Dell Latitude E7450, i7-5600U CPU @ 2.60GHz processor). Fig. 5(a) shows screen shots of the data set used in this section, where the camera moves towards obstacles at different distances. In Fig. 5(b), EdgeFlow scaled by EdgeStereo, now dubbed as Edge-FS, results in the velocity estimates.

Edge-FS is contrasted against the well-known optical flow method developed by Färneback [7], a dense optical flow method (Fig. 5). Although less used than a more conventional KLT-tracker [24], preliminary analyses indicated it to be more suited for the low-resolution images used here. With its default parameters set as in MATLAB R2015b, the sparse magenta line illustrate that the KLT-tracker indeed has difficulties with the low-quality, low-resolution images (128 x 96 pixels).

For Färneback, depth is determined by matching the stereo-images with each other and converting the resulting pixel disparity to a distance. To get velocity, the same line-fit is used as for EdgeFlow[1], but here the whole image is considered rather than the compressed form like the edge histograms. After comparison of the methods with different parameters, both Edge-FS and Färneback are set up with a window size of 11 pixels and a search range of 15 pixels (Färneback's pyramid level at 1). Both forward (x) and sideways (y) velocity measurements, shown in Fig. 5, are compared against the "ground truth" as obtained with an OptiTrack motion capture system, with 24 infrared cameras[2]. The plots also include several values to determine the quality of the velocity estimates: Mean Squared Error (MSE), Variance (VAR) and Normalized Maximum Cross-Correlation Magnitude (NMXM). A low MSE indicates greater similarity and low VAR is a smaller spread of the measurement from the ground truth. A high NMXM stands for a better shape correlation between

[1] The EdgeFlow code as embedded on the stereo camera has a mean computation time for EdgeFlow is 0.00134 seconds (compiled for Linux) and for Färneback is 0.00466 seconds on the same stereo-image data set.
[2] www.optitrack.com

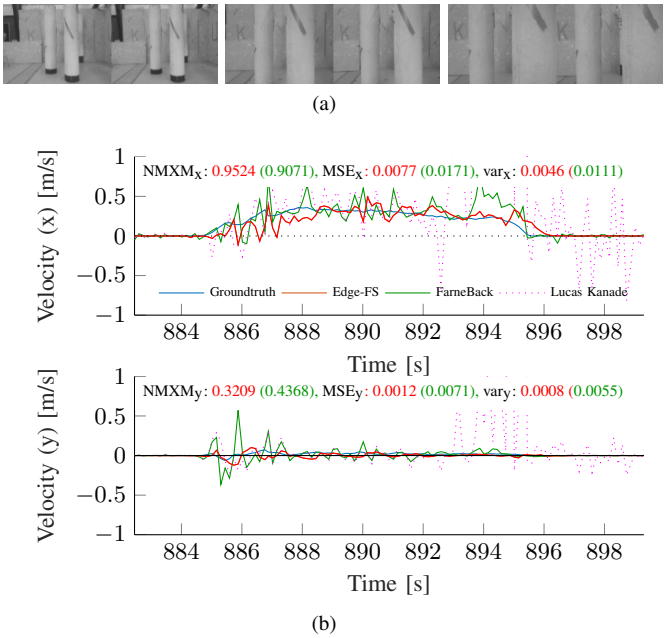

Fig. 5. (a) Several screen shots of the set of images used for off-line estimation of the velocity. Here the diversity in amount of texture can be seen. (b) Off-line velocity estimate calculated by Edge-FS and Färneback, held against the ground truth for the forward moving camera's data set. The quality values are shown for both. The method of Lucas Kanade is shown here but will not be used for the comparison.

the two. All these metrics indicate Edge-FS to obtain more accurate results on this data set than the computationally more expensive Färneback method.

It is important to note that because of the nonlinear relation between pixel disparity and depth in stereo vision, far distances are measured less accurately. The disparities for further distances will become sub-pixel and hard to determine. This is especially relevant to the small stereo board used in this study, which we set up to use 128 x 96 pixel images for the 57.4 x 44.5 deg field of view (FOV). Also the translational optical flow of objects is harder to measure when they are further away, since it becomes sub-pixel as well. Hence, both terms on the left in (3), p_x and d become less accurate at far distances. This correlation between distance and accuracy is shown seen in the box plot of Fig. 6.

Besides the difficulty with larger distances, which is fundamental to stereo vision, Edge-FS also has a bit difficulty determining the forward flow when there is a large lateral motion. Sideways velocity has a 0th order effect on the flow, while forward velocity information is captured by the divergence of the flow field, which is a 1st order effect. Because of this, the forward velocity is more subject to noise and harder to estimate. Later in the experiments, the MAV will mostly fly forward and compensate for sideways drift, so this effect can be countered with filtering and/or a larger window size. Therefore, we will not go into it any further, however we encourage the usage of our MATLAB code for EdgeFlow, with the larger, more diverse, image data-set[3].

[3] https://github.com/knmcguire/EdgeFlow_matlab

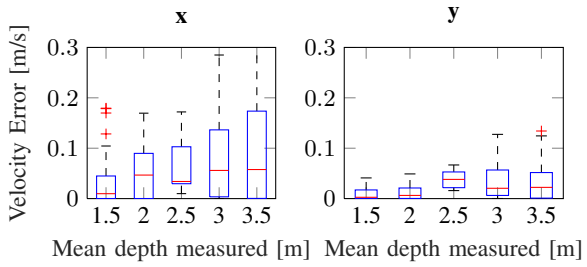

Fig. 6. Boxplot of the absolute velocity estimate error of Edge-FS, compared against the mean observed depth.

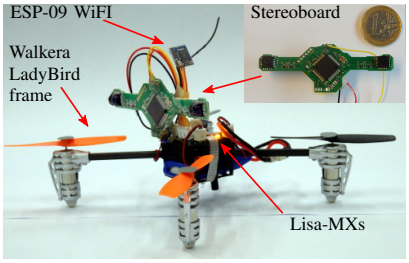

Fig. 7. The 4 g stereo camera mounted on the pocket drone.

III. Experiments on the Pocket Drone

In this section, we explain the implementation of Edge-FS on-board a pocket drone and how it is used in an autonomous obstacle avoidance task. We will first present the velocity estimates by Edge-FS local during flight on the stereo board. Subsequently, a closed loop control is shown, where the drone autonomously navigates through a room, while maintaining its velocity and avoiding obstacles.

A. Hardware specifics

Edge-FS local runs embedded on the stereo camera (as introduced in [4]). Fig. 7 displays two cameras with 1/6 inch image sensors, with a baseline of 6 cm and a Field of View (FOV) of 57.5^o x 44.5^o. The stereo camera has an embedded microprocessor, a STM32F4 with a speed of 168 MHz and 182 kB of memory. The used stereo-images for the on-board algorithm, have a size of 128 x 96 pixels. The maximum reachable frame rate of the stereo camera is 25 Hz, brought down to just 23 Hz with EdgeFlow together with EdgeStereo.

For the experiments, the mounted stereo camera faces forward on a pocket drone, which is the only camera present on the platform. A frame of a Walkera QR LadyBug[4] is adopted as a base. An adapted smaller variant of the Lisa-MX[5] will be used the auto-pilot. The Lisa-MXs also carries a STM32F4 microprocessor, with a speed of 168 Hz and 1 MB of flash memory.ÂăÂă With an ESP-09 WiFi module, telemetry can be broadcasted to the computer to receive all the measured variables required for validation. The entire assembly, including stereo camera and battery, weights exactly 41.9 g.

The auto-pilot program flashed on the Lisa-MXs is Paparazzi[6]. The software runs entirely on-board the microprocessor which governs all the basic flight controls. An Adaptive Incremental Nonlinear Dynamic Inversion (INDI) controller [25] is used for the attitude stabilization of the MAV. The guidance controller resides on top of the stabilization control, to calculate the desired pitch and roll angle, to achieve a desired altitude position or airspeed. In this paper, it will be applied to maintain a desired velocity. It will need the measurements from the stereo camera, operating in parallel with the Lisa-MXs.

B. Velocity Estimate

We have demonstrated in the previous section that Edge-Flow can measure its velocity based on a collection of images. Now it is implemented in the 4 g stereoboard, fixed on a pocket drone. The question is, if it still can see motion, together with all the additional effects caused by motion and vibrations during flight.

Fig. 8(a), presents the velocity estimates of Edge-FS, during a manually controlled flight in front of a textured screen (screen shot in Fig. 8(b) and position in Fig. 8(b)). Its actual velocity is monitored by the same OptiTrack system used for the image data set (Fig. 5(a)). The raw unfiltered velocity measurements of Edge-FS is compared with this ground truth with NMXM, VAR and MSE. Noticeable is that the forward velocity shows more noise peaks than the sideways velocity as expected (see Section II-C). However, in both directions, Edge-FS matches the ground truth adequately, which could be applied for closed-loop control.

To use the actual raw measurements in flight is undesirable. The most common way is to fuse these vision-based velocity estimates with the accelerometers. On a larger MAV than the pocket drone, this would be possible because of the damping. However, because the many vibrations caused by the small propellers being so close to the autopilot, the accelerometers readings contain too much noise. Therefore, in this paper, we only use a median filter with the 5 last measurements, to keep the delay to a minimum.

C. Autonomous Obstacle Avoidance

In the previous subsection we showed validation of the velocity estimate as calculated by Edge-FS. Now we will present a closed-loop flight, where the pocket drone is guided by the velocity estimates, to avoid obstacles identified by means of its stereo vision. The main goal of this experiment is to show the potential of the proposed algorithms to be used for full autonomous navigation. In this section, the vertical position as measured by OptiTrack is exclusively used for height control, as no position measurement is used in the horizontal plane (solely for validation afterwards). This is where the MAV uses its velocity estimates by Edge-FS.

Fig. 9(b) displays the basic control scheme for the navigation task. It determines a desired velocity to avoid collision. The error between the estimate and the desired velocity is the

[4] http://www.walkera.com/ [5] http://wiki.paparazziuav.org/wiki/Lisa/MX [6] http://wiki.paparazziuav.org/

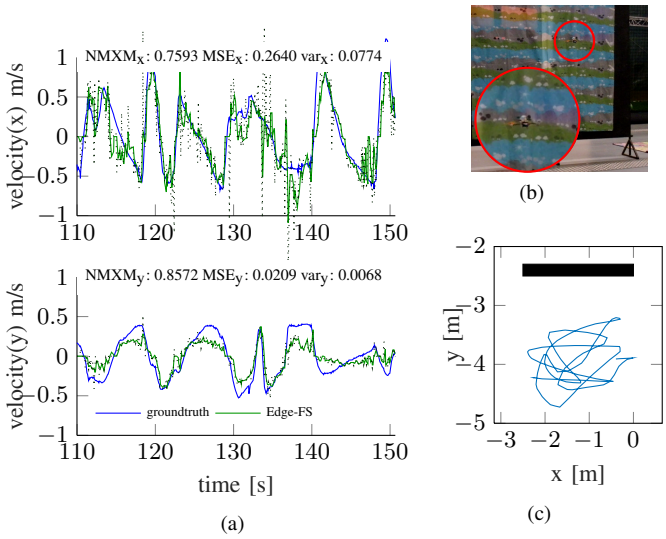

Fig. 8. (a) Velocity estimation by Edge-FS on the pocket drone, (b) a screen shot from the flight in front of a wall and (c) the position during a remote controlled controlled maneuver. Dotted line is the unfiltered Edge-FS.

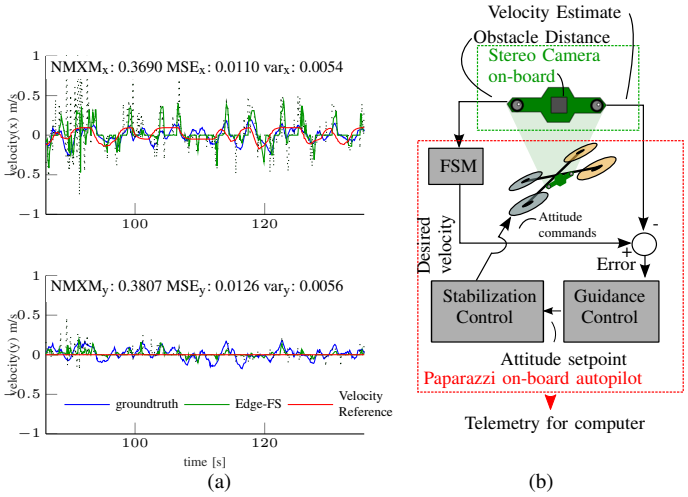

Fig. 9. (a) Velocity control on pocket drone with a wall forcefield and (b) the control scheme used, explaining the hierarchy of the on-board sensors and controllers.

input to the velocity guidance controller, which sets an attitude set-point for the stabilization. Subsequently, EdgeStereo determined the nearest object to camera. If too close, it will produce a backward velocity reference to the guidance controller, 'force field', therefore preventing the pocket drone from hitting the wall face-on. Fig. 9(a) shows the readings from a short flight, where the velocity estimates are used for a simple hover with the obstacle force field.

When encountering a wall/obstacle, the pocket drone will need to move away from the situation. The avoidance scheme is expressed as a simple finite state machine (FSM) with 4 behavioral states (see Fig. 10(b)). It starts in *check* mode, where the pocket drone will check if an obstacle is detected within 1 meters of its FOV by EdgeStereo. If the way is clear, the pocket drone moves *forward* with a constant speed (set now to 0.3 m/s), guided by the velocity estimate from Edge-FS. If it detects an object on its path, the MAV will first hover

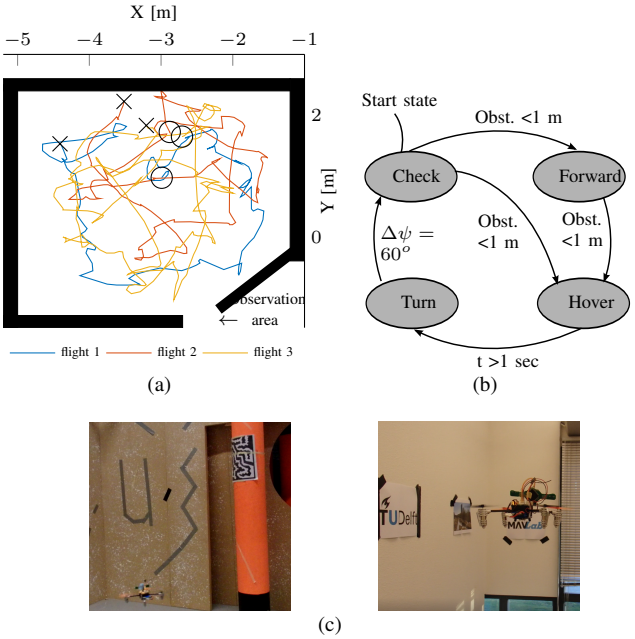

Fig. 10. (a) A position plot of 3 flights, from which the first lasted 91 seconds, the second 101 seconds and the third 122 seconds. (b) shows screen shots of the experiments.

for 1 second and turn with a constant angle relative from the heading (here $\Delta\psi = 60^o$). Immediately after the *check* mode, the MAV will evaluate the situation and proceeds from there.

We performed multiple autonomous flights with the pocket drone. Fig 10(a) shows the result of 3 representative flights of the pocket drone with the forward looking stereo camera. The pocket drone has to navigate in a small room of 4 x 4 meter with varying textured surfaces, (screen shot of camera footage [7] in Fig. 10(c)). All the flights lasted longer than 90 seconds, from which the longest duration was 122 seconds (flight 3). When the pocket drone brushed against the wall, the safety pilot took over the flight with a remote control for a safe landing. The most common case for most of the flights, is that the MAV will approach a wall with a small angle. After the turn, the drone will fly almost parallel to the wall, which it can not detect due to its limited FOV. This is the case for flight 2 and 3, except for flight 1, in which case the pocket drone was facing the observer after a turn. Several flights of the pocket drone has been done within a real-world environment (Fig. 10(c)), which can be observed with the accompanying video and YouTube list.

The failure case mentioned in the flights, is a difficult one to overcome. If the drone would turn, and is then facing an open space, the distance for EdgeStereo could be far enough to compromise the quality for the velocity estimate, when scaled with EdgeFlow. As we already observed in Fig. 6), this would cause the pocket drone to drift significantly when near a wall/obstacle after the turn. If an obstacle is not in its FOV, there is a high chance of collision. This could be solved by merging the *check* and *turn* node of the FSM, so it will

[7] https://www.youtube.com/playlist?list=PL_KSX9GOn2P812tmddfrTlURHNieRe6YY

only stop turning at a significant clear path. Another solution is to add a light weight short range sensor on the sides of the pocket drone, so it will detect at the last minute if the drone is flying aside an obstacle.

The obstacle avoidance logic will need some additional work, however the experiments show that Edge-FS can certainly be used in navigation overall. During the autonomous flight, one could see that the pocket drone was stabilizing itself using the velocity estimates of its forward camera alone.

IV. Conclusion

A computationally efficient optical flow and stereo algorithm is presented in this paper, called Edge-FS. It runs embedded on a very lightweight stereo camera and can be carried by a 40 g pocket drone for determining velocity and depth. In this paper, the forward facing stereo camera is mounted on a pocket drone for an automated control task. It used the velocity estimates to follow a certain desired velocity pattern (given externally), while avoiding obstacles by stereo. The velocity estimate assuming a 3D depth structure shows an improvement in the measurement, however this quality decreases with the distance, which has been taken into account for remainder of the experiments.

In the closed loop flight, the pocket drone with the stereo camera autonomously navigated and avoided obstacles in an area of 4 x 4 meters. Experiments with a simple finite state machine controller have shown that the velocity estimates and the depth measurement can be used for fully autonomous flight. The current work lays the basis for stabilization and collision avoidance on pocket drones with a single, small stereo vision system.

Acknowledgments

We would like to show our gratitude to Kirk Scheper from the MAVlab, Delft University of Technology, for assistance in the development of the guided flight mode and the finite state machine. These results greatly improved the quality of the paper.

References

[1] X. Hu and P. Mordohai, "Evaluation of stereo confidence indoors and outdoors," in *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*. IEEE, 2010, pp. 1466–1473.

[2] H. Hirschmuller, "Stereo processing by semiglobal matching and mutual information," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 30, no. 2, pp. 328–341, 2008.

[3] A. Geiger, J. Ziegler, and C. Stiller, "Stereoscan: Dense 3d reconstruction in real-time," in *Intelligent Vehicles Symposium (IV), 2011 IEEE*. IEEE, 2011, pp. 963–968.

[4] C. De Wagter, S. Tijmons, B. Remes, and G. de Croon, "Autonomous flight of a 20-gram flapping wing mav with a 4-gram onboard stereo vision system," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. IEEE, 2014, pp. 4982–4987.

[5] S. Tijmons, G. de Croon, B. Remes, C. De Wagter, and M. Mulder, "Obstacle avoidance strategy using onboard stereo vision on a flapping wing mav," *arXiv preprint arXiv:1604.00833*, 2016.

[6] B. K. Horn and B. G. Schunck, "Determining optical flow," pp. 185–203, 1981.

[7] G. Farnebäck, "Two-frame motion estimation based on polynomial expansion," in *Scandinavian conference on Image analysis*. Springer, 2003, pp. 363–370.

[8] J. Shi and C. Tomasi, "Good Features To Track," 1994.

[9] E. Rosten and T. Drummond, "Fusing points and lines for high-performance tracking," *International Conference on Computer Vision (ICCV)*, 2005.

[10] D. Honegger, L. Meier, P. Tanskanen, and M. Pollefeys, "An open source and open hardware embedded metric optical flow cmos camera for indoor and outdoor applications," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*. IEEE, 2013, pp. 1736–1741.

[11] D.-J. Lee, R. W. Beard, P. C. Merrell, and P. Zhan, "See and avoidance behaviors for autonomous navigation," *Proceedings of SPIE - The International Society for Optical Engineering*, vol. 5609, pp. 23–34, 2004.

[12] K. McGuire, G. de Croon, C. de Wagter, B. Remes, K. Tuyls, and H. Kappen, "Local histogram matching for efficient optical flow computation applied to velocity estimation on pocket drones," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, May 2016, pp. 3255–3260.

[13] V. Grabe, H. H. Bulthoff, D. Scaramuzza, and P. R. Giordano, "Nonlinear ego-motion estimation from optical flow for online control of a quadrotor UAV," *The International Journal of Robotics Research*, vol. 34, no. 8, pp. 1114–1135, 2015. [Online]. Available: http://ijr.sagepub.com/cgi/doi/10.1177/0278364915578646

[14] H. Romero, S. Salazar, and R. Lozano, "Real-time stabilization of an eight-rotor UAV using optical flow," *IEEE Transactions on Robotics*, vol. 25, no. 4, pp. 809–817, 2009.

[15] F. Kendoul, I. Fantoni, and K. Nonami, "Optic flow-based vision system for autonomous 3d localization and control of small aerial vehicles," *Robotics and Autonomous Systems*, vol. 57, no. 6, pp. 591–602, 2009.

[16] O. Dunkley, J. Engel, J. Sturm, and D. Cremers, "Visual-inertial navigation for a camera-equipped 25g nano-quadrotor," in *IROS2014 aerial open source robotics workshop*, 2014.

[17] F. Ruffier, S. Viollet, S. Amic, and N. Franceschini, "Bio-inspired optical flow circuits for the visual guidance of micro air vehicles," *Proceedings of the 2003 International Symposium on Circuits and Systems, 2003. ISCAS '03.*, vol. 3, pp. 846–849, 2003.

[18] W. E. Green and P. Y. Oh, "Optic-flow-based collision avoidance," *IEEE Robotics & Automation Magazine*, vol. 15, no. 1, pp. 96–103, 2008.

[19] A. Briod, J.-c. Zufferey, and D. Floreano, "Optic-Flow Based Control of a 46g Quadrotor," *IEEE/RSJ International Conference on Robotics and Automation*, 2013.

[20] A. Briod, J.-C. Zufferey, and D. Floreano, "A method for ego-motion estimation in micro-hovering platforms flying in very cluttered environments," *Autonomous Robots*, vol. 40, no. 5, pp. 789–803, 2016.

[21] R. J. D. Moore, K. Dantu, G. L. Barrows, and R. Nagpal, "Autonomous MAV guidance with a lightweight omnidirectional vision sensor," *IEEE International Conference on Robotic Applications*, 2014.

[22] H. Oleynikova, D. Honegger, and M. Pollefeys, "Reactive avoidance using embedded stereo vision for mav flight," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 50–56.

[23] H. C. Longuet-Higgins and K. Prazdny, "The interpretation of a moving retinal image," *Proceedings of the Royal Society of London B: Biological Sciences*, vol. 208, no. 1173, pp. 385–397, 1980.

[24] J.-Y. Bouguet, "Pyramidal implementation of the affine lucas kanade feature tracker description of the algorithm," *Intel Corporation*, vol. 5, no. 1-10, p. 4, 2001.

[25] E. J. Smeur, Q. Chu, and G. C. de Croon, "Adaptive incremental nonlinear dynamic inversion for attitude control of micro air vehicles," *Journal of Guidance, Control, and Dynamics*, vol. 38, no. 12, pp. 450–461, 2015.